REMARKS

In response to the Office Action mailed January 16, 2001, Applicant respectfully requests reconsideration.

Claims 1-32 have been examined. By this amendment, Applicant is amending claims 1, 3, 15, 17, 21, and 26. As a result, claims 1-32 are pending with claims 1, 15, and 21 being independent claims. No new matter has been added.

Claim Rejections - 35 U.S.C. §102(b)

Claims 1-2, 9, 15, and 21-25 were rejected under 35 U.S.C. 102(b) as being anticipated by Sutton, et al. "Processors Sharing Partitioning of Main Storage in the MP System" - IBM TBD v. 22, n. 5, October 1979 (hereinafter "Sutton"). In response, Applicant has amended the claims to distinguish over Sutton.

Sutton is directed to a system for sharing a main memory area in a multi-processor system, wherein the memory is partitioned, such that each processor has its own unique storage space and addressability within memory. In particular, a storage configuration control array (BCA) in the plural processor shared system includes a processor identification field (PRID) (please see page 1 of Sutton). The BCA is tested on each processor access request to the memory address range corresponding to a row in the BCA selected by the request address. Access is permitted or denied to each processor request according to the settings in the PRID fields in the selected row of the BCA. In particular, a one-bit setting in the PRID field assigned to a particular processor will allow that processor to access the range of memory addresses corresponding to that row in the table. The selected row selects the basic storage module (BSM), or part of the BSM in which the absolute storage address range defining the row is physically located. If the bit in the PRID field of the requesting processor is zero, memory selection is inhibited for that processor request.

It is believed that the system implementing Sutton may be shown in more detail in Figure 6 of U.S. Patent Number 4,280,176 issued to R.G. Tan (an author of Sutton), a copy of which is provided for consideration. As shown, the shared memory system is implemented in a system having a number of BSMs connected by a storage controller (SC) to multiple central processors (CPs) 14A and 14B.

In summary, Sutton discloses a multiple processor system having a shared main memory area, analogous to the memory 42 shown in Fig. 3 of the present application.

Claim 1 recites a data management method for managing access to a storage system by at least two devices coupled to the storage system through a network. The method comprises steps of receiving over the network at the storage system a request from one of the at least two devices to access a portion of data stored at the storage system, and selectively servicing, at the storage system, the request responsive to configuration data indicating that the one of the at least two devices is authorized to access the portion of data.

Sutton does not anticipate that which is recited in claim 1. Sutton teaches a single multiprocessor system having a number of basic storage modules (BSMs) that are accessed by multiple central processors (CPs) through a storage controller (SC), but the storage modules do not receive requests over a network. By contrast, claim 1 recites a data management method for managing access to a storage system by at least two devices coupled to the storage system through a network. A request is received from one of the at least two devices to the storage system over the network, and the storage system selectively services the request at the storage system. Therefore, the system of Sutton does not anticipate the method as recited in claim 1. Thus, claim 1 is patentable over Sutton, and the rejection should be withdrawn. Claims 2-14 depend from claim 1 and are patentable for at least the same reasons.

Independent claim 15 recites a computer-readable medium comprising a first data structure to manage accesses by a plurality of devices to volumes of data at a storage system over a communication network, the first data structure comprising a plurality of records corresponding to the plurality of devices, each of the plurality of records corresponding to one of the plurality of devices and including configuration data identifying which of the volumes of the storage system the one of the plurality of devices is authorized to access.

Sutton does not disclose a computer-readable medium as recited in claim 15. In particular, Sutton does not disclose a first data structure that manages "accesses by a plurality of devices to volumes of data at a storage system over a communication network," as recited in claim 15. As discussed above, Sutton only teaches a multiprocessor system and access control of processors to the main memory of that system; Sutton does not manage accesses by a plurality of devices to volumes of data at a storage system over a communication network. Therefore, claim

15 is patentable over Sutton, and the rejection should be withdrawn. Claims 16-20 depend from claim 15 and are patentable for at least the same reasons.

Independent claim 21 recites a storage system comprising at least one storage device apportioned into a plurality of volumes, a configuration table to store centralized configuration data identifying which of the plurality of devices coupled to a storage system via a network are authorized to access each of the plurality of volumes, and a filter, responsive to the configuration data, to selectively forward to the at least one storage device requests for access to the plurality of volumes received from the plurality of devices.

Sutton does not disclose a storage system as recited in claim 21. Sutton teaches a storage configuration control array (BCA), within a single MP system, that controls access to a shared memory region between a subset of processors. Sutton does not, however, teach or suggest a configuration table to store configuration data identifying which of a plurality of devices coupled to a storage system via a network are authorized to access each of the plurality of volumes, and a filter, responsive to the configuration data, to selectively forward to the at least one storage device requests for access to the plurality of volumes received from the plurality of devices, as recited in claim 21. Sutton teaches a control array within a single multiprocessor system, not a filter that is responsive to configuration data that selectively forwards requests received from a plurality of devices over a network. Because Sutton teaches only a single multiprocessor system and access control of processors to memory of that system, Sutton does not teach a storage system operating over a network in the manner claimed. Thus, the system of Sutton does not disclose the storage system recited in claim 21. Therefore, claim 21 is patentable over Sutton, and the rejection should be withdrawn. Claims 22-31 depend from claim 21 and are patentable for at least the same reasons.

Claim Rejections Under 35 U.S.C. §102(e)

Claims 1-4, 9-11, 13, 14, 15-20, 21-27, 30-32 were rejected under 35 U.S.C. 102(e) as being anticipated by Wolff, U.S. Patent No. 5,999,930 (hereinafter "Wolff"). Applicant respectfully traverses the rejection as outlined below.

Wolff is directed to a system for managing distributed control (e.g., distributed among a client and a storage volume) of a shared storage volume without loss of coherency on those volumes (Please see Abstract, Col. 6, lines 22-26). More particularly, a user, logged into a

B

storage volume from a client, is permitted or denied write access to a storage volume based on information stored in an access control table 206 and volume control table 208 (Col. 7, lines 9-17). Each client executes a control process that allows a user to mount and access volumes, perform writes and manage files (Col. 6, lines 29-33 and Col. 7, lines 9-13).

The access control table 206 as shown in Fig. 3A includes an identifier of a client onto which a user is logged. More particularly, a User_Name identifier indicates the name of a user that is searched upon a volume control table logon and which is used for later access (Col. 11, lines 63-67 [Emphasis added]). The access control table stores the password for the user and an indication that indicates whether a process associated with the user (a "user process") is logged on (Col. 8, lines 55-59). Until a user is logged onto the access control table, any applications on that client will not be able to access any of the physical devices whose device drivers have a lock (Col. 8, lines 60-63).

The volume control table indicates, for a user logged onto a client, whether the user is currently locked out for concurrency purposes or allowed to mount a particular storage volume (Col. 9, lines 15-56). This information is used, for example, to determine which user has current write access to a particular volume, and which users have mount access privileges for each specific volume (Col. 9, lines 32-33 and Col. 9, lines 37-40). There is one volume control table for each volume which supports distributed control by multiple writers (Col. 9, lines 18-20). Information stored in the access and volume access tables is transferred to the client, where associated privileges are applied to the local reference to the remote files stored on the volume (Col. 8, lines 6-15).

In summary, Wolff discloses a user-based process-to-process authorization system that manages accesses by users to volumes on a storage server, wherein the system manages concurrency (e.g., does not allow concurrent writes by two users) and determines whether a user can mount a volume of the storage server.

Wolff does not anticipate that which is recited in claim 1. More particularly, Wolff does not disclose selectively servicing, at the storage system, a request responsive to configuration data indicating that the one of the at least two devices is authorized to access the portion of the data, as recited in claim 1. Wolff authorizes user processes, not devices. Wolff is concerned with authorizing user processes that execute on clients, wherein the identifier used to index a volume control table is a user name having a character string identifier that identifies a user (Col.

B

11, lines 63-67). This user name is used by user processes to access volumes of a storage server. By contrast, claim 1 recites a data management method for managing access to a storage system by at least two devices coupled to the storage system through a network. The method therefore may be implemented to overcome security problems associated with a device misrepresenting its identity to obtain data at the resource (Please see page 3, line 24 through Page 3, line 19 of the instant application). Wolff performs user authorization which is independent of any device, as Wolff is not concerned with authorizing devices; Wolff is concerned with allowing multiple users to write to a volume while maintaining concurrency, and with authorizing users to mount particular volumes based on their user credentials (Col. 9, lines 32-33 and Col. 9, lines 37-40). Therefore, Wolff does not anticipate that which is recited in claim 1 and the rejection should be withdrawn. Claims 2-14 depend from claim 1 and are allowable for at least the same reasons.

Wolff does not anticipate independent claim 15. More particularly, Wolff does not disclose a first data structure to manage accesses by a plurality of devices to volumes of data at a storage system over a communication network, the first data structure corresponding to one of the plurality of devices and including configuration information identifying which of the volumes of the storage system the one of the plurality of devices is authorized to access, as recited in claim 15. As discussed above with respect to claim 1, Wolff authorizes user processes, not devices. Thus, Wolff does not identify which volumes that one of a plurality of devices is authorized to access. Therefore, Wolff does not anticipate that which is recited in independent claim 15, and the rejection should be withdrawn. Claims 16-20 depend from claim 15 and are allowable for at least the same reasons.

Wolff also does not anticipate claim 21. More particularly, Wolf does not disclose a configuration table to store configuration data identifying which of a plurality of devices coupled to a storage system via a network are authorized to access each of the plurality of volumes, as recited in claim 21. As discussed above with respect to claim 1, Wolff authorizes user processes, not devices. Thus, Wolff does not identify devices that are authorized to access a plurality of volumes. Therefore, Wolff does not anticipate that which is recited in independent claim 21, and the rejection should be withdrawn. Claims 22-31 depend from claim 21 and are allowable for at least the same reasons.

Claim Rejections Under 35 U.S.C. §103(a)

Claims 5-8 and 28-29 were rejected under 35 U.S.C. 103(a) as being obvious over Wolff, in view of Russel, U.S. Patent No. 5,455,953 (hereinafter "Russel"). Dependent claims 5-8 and 28-29 are allowable for at least the same reasons as the independent claims from which they depend. Further, Russel does not supply the above-discussed missing limitations from the independent claims, as Russel discloses a user-based authorization; Russel does not authorize devices (Please see Col. 21, line 48 through Col. 23, line 50). Therefore, claims 5-8 and 28-29 are believed to be in allowable condition, and the rejection of claims 5-8 and 28-29 should be withdrawn.

CONCLUSION

In view of the foregoing amendments and remarks, this application should now be in condition for allowance. A notice to this effect is respectfully requested. If the Examiner believes, after this amendment, that the application is not in condition for allowance, the Examiner is requested to call the Applicant's attorney at the telephone number listed below.

If this response is not considered timely filed and if a request for an extension of time is otherwise absent, Applicant hereby requests any necessary extension of time. If there is a fee occasioned by this response, including an extension fee, that is not covered by an enclosed check, please charge any deficiency to Deposit Account No. 23/2825.

Respectfully submitted
S. Blumenau, et al., Applicant(s)

By: _____

Edward J. Russavage, Reg. No. 43,069
Richard F. Giunta, Reg. No. 36,149
Wolf, Greenfield & Sacks, P.C.
600 Atlantic Avenue
Boston, Massachusetts 02210-2211
Tel. No.: (617) 720-3500
Attorney's for Applicant(s)

Docket No. E0295/7066 (RFG/EJR)
Dated: April 16, 2001
x04/16/01

B

AMENDED CLAIMS SHOWING CHANGES

1. (Amended) A data management method for managing access to a storage system by at least two devices coupled to the storage system through a network, the method comprising [a] steps of:

receiving over the network at the storage system a request from one of the at least two devices for access to a portion of data stored at the storage system; and

selectively servicing, at the storage system, [a] the request [from one of the at least two devices for access to a portion of data stored at the storage system] responsive to configuration data indicating that the one of the at least two devices is authorized to access the portion of data.

3. (Amended) The data management method according to claim 1, [wherein the storage system and the at least two devices are coupled together by a network,] the method including a step of:

forwarding the request from the one of the at least two devices to the storage system over the network.

15. (Amended) A computer readable medium comprising:

a first data structure to manage accesses by a plurality of devices to volumes of data at a storage system over a communication network, the first data structure comprising a plurality of records corresponding to the plurality of devices, each [record] of the plurality of records corresponding to one of the plurality of devices and including configuration information identifying which of the volumes of the storage system the one of the plurality of devices is authorized to access.

17. (Amended) The combination according to claim 16, in further combination with the plurality of devices and [a] the communication network, wherein the storage system and the plurality of devices are coupled to communicate over the communication network.

21. (Amended) A storage system comprising:

at least one storage device apportioned into a plurality of volumes;

B

a configuration table to store configuration data identifying which of a plurality of devices coupled to the storage system via a network are authorized to access each of the plurality of volumes; and

a filter, responsive to the configuration data, to selectively forward to the at least one storage device requests for access to the plurality of volumes received from the plurality of devices over the network.

26. (Amended) The storage system according to claim 21, in combination with the plurality of devices and [a] wherein the network [that] couples the storage system to the plurality of devices.